

Appl. No. 09/503,608
Amdt. dated June 21, 2004
Reply to Office Action of April 22, 2004

Remarks

The present amendment responds to the Official Action dated April 22, 2004. The Official Action rejected claims 1-10 and 12-14 under 35 U.S.C. § 103(a) based on Wesinger, Jr. et al. U.S. Patent No. 6,052,788 (Wesinger) in view of Reid et al. U.S. Patent No. 6,182,226 (Reid). Claim 11 was rejected under 35 U.S.C. § 103(a) based on Wesinger and Reid and further in view of Bechtolsheim et al. U.S. Patent No. 6,515,963 (Bechtolsheim). These grounds of rejection are addressed below following a brief discussion of the present invention to provide context.

Claim 1 has been amended to modify the term "datagram" with the term "connectionless" in the discarding step to be consistent with the previous usage in the determining step of the claim. Claims 1-14 are presently pending.

The Present Invention

The present invention recognizes that the consequences of intentional datagram flooding attacks and unintentional overload situations resulting from a burst of connectionless datagrams can be mitigated by dropping the traditional notion of attempting to distinguish between legitimate and illegitimate traffic. In the present invention, both legitimate and illegitimate datagram traffic is subject to a common policy that attempts to guarantee that legitimate work will be performed and a server will not crash in flooding situations, irrespective of whether the flooding is caused by legitimate or illegitimate datagram traffic. The present invention helps to

Appl. No. 09/503,608
Amdt. dated June 21, 2004
Reply to Office Action of April 22, 2004

prevent a server from crashing due to overload and it prevents one or more attackers from consuming all resources on a network server.

According to the present invention, in response to the arrival of a datagram destined for a specified port on a network server, the transmitting host is identified from the datagram and the number of datagrams already queued for the same host and for the same port is determined. If this number exceeds a prescribed threshold, the datagram is discarded.

The prescribed threshold is dynamically determined in a presently preferred embodiment. The owner of the network server specifies for each port that is subject to datagram flooding a maximum number of queued datagrams (M) allowed at any time to the port and a controlling percentage (P) of available queue slots remaining for the port. The present invention keeps track of the number (A) of queued datagrams for the port and it calculates the number of available queue slots (I) by subtracting the number of queued datagrams from the maximum number of datagrams ($I = M - A$). If the number of datagrams already queued for the transmitting host is equal to or greater than P times the number of queue slots left ($M \geq P \cdot I$), then the present datagram is not queued for the port. Otherwise, the datagram is queued and the number of queued datagrams (A) for the port is incremented by one.

The Art Rejections

As addressed in greater detail below, Wesinger, Reid, and Bechtolsheim do not support the Official Action's reading of them and the rejections based thereupon should be reconsidered and withdrawn. Further, the Applicant does not acquiesce in the analysis of the relied upon art

Appl. No. 09/503,608
Amdt. dated June 21, 2004
Reply to Office Action of April 22, 2004

made by the Official Action and respectfully traverses the Official Action's analysis underlying its rejections.

Claims 1-10 and 12-14 were rejected under 35 U.S.C. § 103(a) based on Wesinger in view of Reid. Wesinger describes a firewall which employs envoys which combine the security robustness described in prior-art proxies and the transparency and ease-of-use of prior-art packet filters. Wesinger, Abstract. To achieve full transparency, the firewall is configured as two or more sets of virtual hosts. One set of hosts responds to addresses on a first network interface of the firewall. Another set of hosts responds to addresses on a second network interface of the firewall. Before traffic can pass through the firewall, an envoy must be established for that traffic. Wesinger, col. 3, lines 60-62.

As described at col. 13, lines 27-48, authentication rules checking is performed on a first data packet to be sent from a first computer to a second computer. If the result of this rules checking is to allow the first packet to be sent, a time-out limit associated with communications between the first computer and the second computer via UDP is established, and the first packet is sent from one of the virtual hosts to the second computer on behalf of the first computer. Thereafter, for so long as the time-out limit has not expired, subsequent packets between the first computer and the second computer are checked and sent. After the time-out limit has expired, the virtual host may be remapped to a different network address to handle a different connection. Typical authentication rules include restricting access to a known secure host, and requiring username/password authentication. By allowing traffic to pass before expiration of a timer,

Appl. No. 09/503,608
Amdt. dated June 21, 2004
Reply to Office Action of April 22, 2004

Wesinger's envoy handles connectionless traffic in a totally different manner than the present invention.

In contrast, the present invention addresses a method for defending against network flooding attacks of connectionless datagrams. In particular, the method determines, in response to the arrival of a connectionless datagram from a host for a port on a network server, if the number of connectionless datagrams already queued to the port from the host exceeds a prescribed threshold. If so, the method discards the connectionless datagram which advantageously prevents a particular host from flooding a particular port. If the number of connectionless datagrams already queued to the port from the host does not exceed the prescribed threshold, the connectionless datagram is queued to a queue slot of the port. Claim 1, as presently amended, reads as follows:

1. A method of preventing a flooding attack on a network server in which a large number of connectionless datagrams are received for queuing to a port on the network server, comprising:
determining, in response to the arrival of a connectionless datagram from a host for a port on the network server, if the number of connectionless datagrams already queued to the port from the host exceeds a prescribed threshold;
discarding the connectionless datagram, if the number of connectionless datagrams already queued to the port from the host exceeds the prescribed threshold; and
queuing the connectionless datagram to a queue slot of the port, if the number of connectionless datagrams already queued to the port from the host does not exceed the prescribed threshold. (emphasis added)

Wesinger does not teach and does not suggest "determining, in response to the arrival of a connectionless datagram from a host for a port on the network server, if the number of connectionless datagrams already queued to the port from the host exceeds a prescribed

Appl. No. 09/503,608
Amdt. dated June 21, 2004
Reply to Office Action of April 22, 2004

threshold," as presently claimed. Further, Wesinger does not teach and does not suggest "discarding the connectionless datagram, if the number of connectionless datagrams already queued to the port from the host exceeds the prescribed threshold," as presently claimed. Wesinger merely provides authentication rules in combination with a timer to allow traffic to pass through the firewall.

The Official Action relies on the text found at col. 14, lines 22-31, col. 14, lines 36-37, and col. 7, lines 1-4 of Wesinger as purportedly suggesting the determining, discarding, and queuing steps in claim 1. Applicants respectfully disagree. These relied upon portions of text address connection oriented protocols and not connectionless datagrams as claimed. It should be noted that the disclosure at col. 13, lines 27-48 of Wesinger which is discussed above addresses Wesinger's approach to connectionless datagrams. That approach is clearly different than that which is presently claimed.

Reid fails to cure the deficiencies of Wesinger as a reference. Reid describes a firewall used to achieve network separation within a computing system having a plurality of network interfaces. A plurality of regions is defined within the firewall and a set of policies is configured for each of the plurality of regions where each network interface is assigned to only one region. Reid, Abstract and col. 1 lines 64-65. Reid's firewall restricts communication to and from each of the plurality of network interfaces in accordance with the set of policies configured for the region assigned to the network interfaces carrying the communication. Reid, col. 1, line 67 - col. 2, line 4. To program the set of policies, Reid utilizes an access control language. Such access control language allows policies to restrict access to communication by utilizing criteria such as

Appl. No. 09/503,608
Amdt. dated June 21, 2004
Reply to Office Action of April 22, 2004

source and destination region, users and groups, load balancing, and maximum number of concurrent sessions. Reid, col. 7 lines 39-58. These varied criteria of Reid do not make obvious the presently claimed approaches of preventing a flooding attack on a network server.

The Official Action apparently relies on Reid solely for the end result of "a method of preventing a flooding attack on a network server." Although Reid's disclosure states that its system is designed to defend against known network penetration and denial of service attacks such as a SYN flood attack, Reid's specific approach of programming policies to specific network regions utilizing access rules to limit communication access is quite different than the present invention. Reid's disclosure is silent with respect to specific steps to defend against such an attack.

Combining Reid and Wesinger as the Official Action suggests would still fall short of meeting the presently claimed features. Reid and Wesinger, separately or in combination, do not teach and do not suggest "determining, in response to the arrival of a connectionless datagram from a host for a port on the network server, if the number of connectionless datagrams already queued to the port from the host exceeds a prescribed threshold," as presently claimed in claim 1. Reid and Wesinger, separately or in combination, do not teach and do not suggest "discarding the connectionless datagram, if the number of connectionless datagrams already queued to the port from the host exceeds the prescribed threshold," as presently claimed. See also claims 3, 5, and 7.

Dependent claim 11 was rejected under 35 U.S.C. § 103(a) based on Wesinger and Reid and further in view of Bechtolsheim. Bechtolsheim fails to cure the deficiencies of Wesinger and

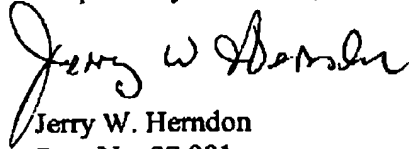
Appl. No. 09/503,608
Amdt. dated June 21 2004
Reply to Office Action of April 22 2004

Reid. Since claim 11 depends from and contains all the limitations of claim 1 as presently amended, claim 11 distinguishes from the references in the same manner as claim 1.

Conclusion

All of the presently pending claims, as amended, appearing to define over the applied references, withdrawal of the present rejection and prompt allowance are requested.

Respectfully submitted,



Jerry W. Herndon
Reg. No. 27,901
IBM Corporation
T81/503, 3039 Cornwallis Road
RTP, NC 27709
(919) 543-3754